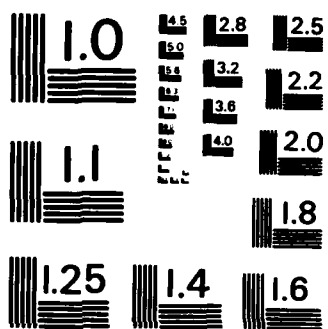AD-A146 025   KNOWLEDGE REPRESENTATION AND NATURAL-LANGUAGE SEMANTICS   1/1
               (U) SRI INTERNATIONAL MENLO PARK CA ARTIFICIAL
               INTELLIGENCE CENTER   R C MOORE 13 JUL 84
UNCLASSIFIED   AFOSR-TR-84-0799 F49620-82-K-0031   F/G 5/7   NL

MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS-1963-A

AFOSR-TR-

KNOWLEDGE REPRESENTATION AND NATURAL-LANGUAGE SEMANTICS

July 13, 1984

2nd Annual Technical Report
Covering the Period June 1, 1983 to May 30, 1984

SRI Project 4488
Contract No. F49620-82-K-0031

By: Robert C. Moore
Staff Scientist

Artificial Intelligence Center
Computer Science and Technology Division

Prepared for:

Air Force Office of Scientific Research
Building 410
Bolling Air Force Base
Washington, D.C. 20332

Attention: Dr. Robert Buchal

DTIC
ELECTE
SEP 26 1984
A

84 09 17 019

I OBJECTIVES OF THE RESEARCH EFFORT

Central to almost all aspects and applications of artificial intelligence is the representation and manipulation of large bodies of knowledge about the world. When viewed from the perspective of their ability to express facts about the external world, however, most knowledge representation schemes currently used in artificial intelligence are constrained by the limits of first-order logic. That is, they provide terms for referring to individuals, predicates for expressing properties and relations of individuals, and mechanisms that achieve some of the effects of propositional connectives and quantifiers. Much research effort has been expended on ways of organizing knowledge bases and developing information retrieval mechanisms; in terms of pure expressive power, however, existing representation systems are rather limited.

This issue is brought into sharp focus when one seriously attempts to analyze *the semantic* content of expressions in natural language, since many types of linguistic expressions seem to require something beyond first-order logic to represent their meaning perspicuously. Specifically, natural languages have special features for dealing with a variety of concepts that are central to our commonsense understanding of the world. For instance, linguistic systems of tense and aspect are intimately connected with commonsense conceptions of time. Adverbial modification, nominalization phenomena, and categorical distinctions among verb phrases appear to depend on such notions as state, event, and process. Predicate complement constructions frequently involve concepts of "propositional attitude" such as knowledge, belief, desire, and intention. The linguistic features of singular/plural and mass/count are used to sort out individuals, collective entities, and substances. In all these cases, either it is not clear how to express these concepts in first-order logic at all—or it is clear that they can be expressed in first-order logic only by very indirect means.

This project undertakes a program of basic research in knowledge representation, focusing on the representation of concepts needed for the semantic analysis of natural language. The objectives of the project are to produce formalisms, suitable for manipulation by computer, for the representation of specific concepts that are important for natural-language semantics, and to give an independent account of the meaning of such representations using the tools of formal logic.

II   STATUS OF THE RESEARCH EFFORT

A.   Representing the Dependence of Action on Knowledge

One of the representational problems we have studied is the relationship between knowledge and action. Both knowledge and action are among the basic concepts that underlie many different areas of commonsense and expert knowledge, but the interaction between the two is particularly important when applying artificial intelligence techniques to planning.

Planning sequences of actions and reasoning about their effects is one of the most thoroughly studied areas within artificial intelligence, but relatively little attention has been paid to the important role that an agent's knowledge plays in planning and acting to achieve a goal. Virtually all planning systems in artificial intelligence are designed to operate with complete knowledge of all relevant aspects of the problem domain and problem situation. Often any statement that cannot be inferred to be true is assumed to be false. In the real world, however, planning and acting must frequently be performed without complete knowledge of the situation.

This constraint imposes two additional burdens on an intelligent agent trying to act effectively. First, when the agent entertains a plan for achieving some goal, he must consider not only whether the physical prerequisites of the plan have been satisfied, but also whether he has all the information necessary to carry out the plan. Second, he

2

...des
ial and/or
Special

must be able to reason about what he can do to obtain necessary information that he lacks. For example, to call someone on the telephone, just being physically able to dial a telephone is not sufficient; one must also know the person's telephone number. One can plan to acquire this information, however, by looking up the number in a telephone book.

During the past year, we have refined and extended our previous work on the dependence of action on knowledge [1]. Our main thesis is that the knowledge required for an action can be analyzed as a matter of knowing what action to take. An agent could know that to call Smith on the telephone he needs to dial Smith's telephone number, but still not know what to do because he does not know precisely what action dialing Smith's telephone number is. That is, he might not know whether dialing Smith's telephone number is the action of dialing 221-1111, or dialing 221-1112, or dialing 221-1113, and so on. We may assume he has a general procedure for dialing telephone numbers, but unless he knows which number to apply it to, he does not, in the relevant sense, know what to do.

In our previous work, we successfully applied this analysis to actions that are treated as nondecomposable wholes, but our treatment of complex plans was less satisfactory. To represent complex plans, we introduced concepts of sequential actions, conditional actions, and iterated actions. Formalizing the knowledge prerequistes of these complex actions was somewhat ad hoc, however. In particular, for conditional actions ("if P is true, then do ACTION1, otherwise do ACTION2") we had to state independently the fact that, in order to carry out a conditional action, an agent must know if the condition is true.

The work performed during the past year remedies this and a number of other deficiencies. The key change is to view a complex plan as a description of a sequence of actions. Then the knowledge prerequisites of complex plans can be given a treatment similar to that for simple actions, so that the agent is assumed to have sufficient knowledge to carry out a plan if he knows what sequence of actions the plan

3

describes. The problem of conditional actions is handled automatically, because what action is described by a conditional action description depends on whether the condition is true. Hence an agent must know whether the condition is true to know what action this is. This work is presented in full in [2].

B. <u>Semantic</u> <u>Analysis</u> <u>of</u> <u>Adverbial</u> <u>Modifiers</u> <u>and</u> <u>Event</u> <u>Sentences</u>

A good example of the way a careful analysis of the meaning of natural-language expressions gives us insight into the representation of commonsense knowledge is presented by our work on the adverbial modification of event sentences. Whether or not there is a fundamental semantic distinction between event sentences, such as "John went to New York," and stative sentences, such as "John was in New York," is one of the more puzzling problems in representing the meaning of expressions in ordinary English. The latter sentence can be analyzed as saying simply that a certain relation, that of location, held between John and New York at some past time. This type of analysis seems less satisfactory, though, for the former sentence. "Went" does not seem merely to express a relation the way "is in" does. Rather, it appears to describe an event, indicated by the fact that it makes sense to ask "When did it happen?" after being told "John went to New York," but not after being told "John was in New York."

One suggestion as to how event sentences might differ from stative sentences is provided by Davidson [3], who suggests that event sentences be represented as explicitly asserting the existence of the event being described. Roughly speaking, this amounts to treating "John went to New York" as if it were "There was a going of John to New York." Davidson's suggestion is intriguing, but, heretofore, there has been relatively little evidence to support it. The study of adverbial modification of event sentences conducted under this project has provided the most convincing support to date for the kind of representation of event sentences given by Davidson and has cleared up several related problems. This work is described more fully in a paper by Croft [4].

4

To summarize this work briefly, we have developed a unified analysis for most "-ly" adverbs and adjectives, namely, as predicates. A small class of adverbs, all indicating modality or uncertainty ("possibly," "probably," "allegedly," etc.), must be treated as modal operators over propositions, as their semantics implies: thus, "John probably ate the cookie" would be represented as PROBABLE[EAT(JOHN,COOKIE)]. The corresponding adjectival forms are interpreted, using restricted quantification notation, as modal operators over the description; thus, "any possible solution" will be (ANY X: POSSIBLE[SOLUTION(X)]).

All other adjectives and adverbs that have the property of "factivity" (viz., if the sentence with the adverb/adjective is true, then the sentence without the adverb/adjective is also true), are predicates. The presence of " -ly" is syntactically determined: if the predicate is modifying a verb or adjective instead of a noun, the "-ly" is added. The semantic difference between "adjectives" and "adverbs" is that the former are the properties of objects, the latter of events, events being represented as event variables following Davidson [3]. Thus, "John slowly entered the room" is ENTER(E,JOHN,ROOM) & SLOW(E).

There are two unusual cases, which must be accounted for. First, a sentence like "Maggie rudely spoke to the Queen" is ambiguous between a manner reading ("The manner in which Maggie spoke to the Queen was rude") and a fact reading ("The fact that Maggie spoke to the Queen was rude"). While the first reading is represented by modification of the event variable, the second reading represents an assertion about a state of affairs, the state of affairs of the proposition "Maggie spoke to the Queen" being true, which we represent by the FACT operator. Thus the two readings are SPEAK(E,MAGGIE,QUEEN) & RUDE(E) and SPEAK(E,MAGGIE,QUEEN) & RUDE(FACT[SPEAK(E,MAGGIE,QUEEN)]) respectively. Second, adverbs of intention ("intentionally," "willingly," etc.), which display referential opacity and other intensional behavior, must be represented as predicates taking an agent and a proposition as well as an event.

All possible derivational patterns between adverbs and adjectives are found. Adverbs like "bitterly," which take an individual and an event, are derived from adjectives that take an individual and describe his emotional state. Adverbs like "slowly," which take an event only, have derived adjectives that take an individual and a role: "John ran the mile fast" vs. "John is fast (at running the mile)." Finally, for adverbs like "rudely" or "cleverly," which take an individual and an event (or FACT operator), the corresponding adjectives are identical in semantic form: in the manner reading, "John cleverly solved the problem" and "John was clever at solving the problem" are both represented as SOLVE(E, JOHN,PROBLEM) & CLEVER(E).

Adjectives and adverbs that are "gradable" (viz., can be modified by degree terms or placed in comparative constructions) will have additional arguments in the predicate structure, and that is being investigated in other work on this project. The fact that gradability applies to both adjectives and adverbs, however, is another confirmation of their underlying semantic unity.

C. New Results on Autoepistemic Logic

In our previous work [5] we developed a nonmonotonic logic for modeling the beliefs of ideally rational agents who reflect on their own beliefs. We called this system "autoepistemic logic." We defined a simple and intuitive semantics for autoepistemic logic, and we were able to show that the logic was both sound and complete with respect to this semantics. However, the nonconstructive character of both the logic and its semantics made it difficult to prove the existence of sets of beliefs satisfying all the constraints of autoepistemic logic. We have recently developed an alternative, possible-world semantics for autoepistemic logic that enables us to construct finite models for autoepistemic theories and to demonstrate the existence of sound and complete autoepistemic theories based on given sets of premises. This work is presented in [6].

6

D. The Deduction Model of Belief

Reasoning about the knowledge and beliefs of computer and human agents is assuming increasing importance in artificial intelligence systems for natural-language understanding, planning, and knowledge representation. A natural model of belief for robot agents is the deduction model: an agent is represented as having an initial set of beliefs about the world in some internal language and a deduction process for deriving some (but not necessarily all) logical consequences of these beliefs. Because the deduction model is an explicitly computational model, it is possible to take into account limitations of an agent's resources when reasoning.

This project has provided partial support for an investigation of a Gentzen-type formalization of the deductive model of belief. Several original results have been proven. Among these are soundness and completeness theorems for a deductive belief logic; a correspondence result that relates our deduction model to competing possible-world models; and a modal analog to Herbrand's Theorem for the belief logic. Specialized techniques for automatic deduction based on resolution have been developed using this theorem.

Several other topics of knowledge and belief have been explored from the viewpoint of the deduction model, including a theory of introspection about self-beliefs, and a theory of circumscriptive ignorance, in which facts an agent doesn't know are formalized by limiting or circumscribing the information available to him. These results are presented in the Ph.D. dissertation of Konolige [7] and are summarized in a shorter paper [8].

III   PUBLICATIONS

Robert C. Moore, "Semantical Considerations on Nonmonotonic Logic," accepted for publication in *Artificial Intelligence*.

Robert C. Moore, "A Formal Theory of Knowledge and Action," to appear in *Formal Theories of the Commonsense World*, J. Hobbs and R. C. Moore, eds. (Ablex Publishing Corporation, Norwood, New Jersey, 1984).

Kurt Konolige, "Belief and Incompleteness," to appear in *Formal Theories of the Commonsense World*, J. Hobbs and R. C. Moore, eds. (Ablex Publishing Corporation, Norwood, New Jersey, 1984).

In preparation:

Robert C. Moore, "Possible-World Semantics for Autoepistemic Logic."

William Croft, "Issues in the Logical Form of Adverbs."

IV   CONFERENCE PRESENTATIONS

Robert C. Moore, "Semantical Considerations on Nonmonotonic Logic," Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, August 8-12, 1983.

V   PERSONNEL

The research of this project has been carried out during the past year by Robert C. Moore and Kurt Konolige, with William Croft as a graduate research assistant. Supervision has been provided by Nils Nilsson and Stanley Rosenschein.

Advanced degrees awarded:

Kurt Konolige, Ph.D., Department of Computer Science, Stanford University, June 1984, dissertation title: *A Deduction Model of Belief and its Logics* (partial support provided by this project).

8

REFERENCES

1. R. C. Moore, "Reasoning about Knowledge and Action," SRI Artificial Intelligence Center Technical Note 191, SRI International, Menlo Park, California (October 1980).

2. R. C. Moore, "A Formal Theory of Knowledge and Action," SRI Artificial Intelligence Center Technical Note 320, SRI International, Menlo Park, California (February 1984).

3. D. Davidson, "The Logical Form of Action Sentences," in The Logic of Decision and Action, N. Rescher, ed., pp. 81-95 (University of Pittsburgh Press, Pittsburgh, Pennsylvania, 1967).

4. W. Croft, "Issues in the Logical Form of Adverbs," in preparation.

5. R. C. Moore, "Semantical Considerations on Nonmonotonic Logic," SRI Artificial Intelligence Center Technical Note 284, SRI International, Menlo Park, California (June 1983).

6. R. C. Moore, "Possible-World Semantics for Autoepistemic Logic," in preparation.

7. K. Konolige, A Deduction Model of Belief and its Logics, Ph.D. dissertation, Department of Computer Science, Stanford University (June 1984).

8. K. Konolige, "Belief and Incompleteness," SRI Artificial Intelligence Center Technical Note 319, SRI International, Menlo Park, California (January 1984).

END

FILMED

10-84

DTIC